

**IN THE CLAIMS:**

1           1. (Original) A system adapted to correct multiple storage device failures in a  
2 storage array using a combination of multiple first parity groups and a single secondary  
3 parity group, the system comprising:  
4           a storage array having a plurality of concatenated sub-arrays, each sub-array in-  
5 cluding a set of data storage devices and a first parity storage device, the array further  
6 including a global secondary storage device associated with the storage array and holding  
7 secondary parity values for the single secondary parity group, the secondary parity values  
8 computed across the concatenation of the sub-arrays.

1           2. (Original) The system of Claim 1 wherein the first parity group is a row parity  
2 group, the first parity storage device is a row parity storage device and wherein each row  
3 parity group is associated with a sub-array of the storage array such that the array is com-  
4 posed of the multiple row parity groups.

1           3. (Original) The system of Claim 2 wherein the secondary parity group is a di-  
2 agonal parity group, the secondary storage device is a diagonal parity storage device and  
3 wherein the secondary parity values are diagonal parity values.

1           4. (Original) The system of Claim 3 further comprising:  
2           a storage operating system configured to implement double failure protection en-  
3 coding of the concatenated sub-arrays, wherein row parity values for each sub-array are

4 stored on each row parity storage device and diagonal parity values for the entire array  
5 are stored on the global diagonal parity storage device; and  
6 a processing element configured to execute the storage operating system to  
7 thereby invoke storage access operations to and from the array in accordance with the  
8 double failure protection encoding.

1 5. (Original) The system of Claim 4 wherein the double failure protection encod-  
2 ing is row-diagonal parity encoding.

1 6. (Original) The system of Claim 4 wherein the double failure protection en-  
2 coding is EVENODD parity encoding.

1 7. (Original) The system of Claim 1 wherein each sub-array is organized as a  
2 concentrated parity disk array.

1 8. (Original) The system of Claim 1 wherein each sub-array is organized as a  
2 distributed parity disk array.

1 9. (Original) The system of Claim 1 wherein the storage devices are video tape,  
2 magnetic tape, optical, DVD, bubble memory, electronic random access memory or mag-  
3 netic disk devices.

1           10. (Original) A method for correcting double failures in a storage array using a  
2 combination of a single diagonal parity group and multiple row parity groups, the method  
3 comprising the steps of:

4           organizing the storage array as a plurality of concatenated sub-arrays based on  
5 double failure protection encoding, each sub-array including a set of data storage devices  
6 and a row parity storage device, the storage array further including a global diagonal par-  
7 ity storage device for holding diagonal parity;

8           computing the diagonal parity for the single diagonal parity group across the con-  
9 catenated sub-arrays; and

10          correcting storage device failure within the array using the row parity storage de-  
11 vice associated with each sub-array and the global diagonal parity storage device associ-  
12 ated with the storage array.

1           11. (Previously presented) A method for correcting double failures in a storage  
2 array using a combination of a single diagonal parity group and multiple row parity  
3 groups, the method comprising the steps of:

4           organizing the storage array as a plurality of concatenated sub-arrays based on  
5 double failure protection encoding, each sub-array including a set of data storage devices  
6 and a row parity storage device, the storage array further including a global diagonal par-  
7 ity storage device for holding diagonal parity;

8           computing the diagonal parity for the single diagonal parity group across the con-  
9 catenated sub-arrays;

10          correcting storage device failure within the array using the row parity storage de-  
11 vice associated with each sub-array and the global diagonal parity storage device associ-  
12 ated with the storage array;

13          encoding the double failure protection as row-diagonal parity encoding;

14           determining whether the storage device failure is to a single storage device in a  
15 sub-array;

16           if the storage device failure is to a single storage device in the sub-array, recon-  
17 structing the failed storage device using local row parity associated with the sub-array;  
18 and

19           if the storage device failure is not to a single storage device in the sub-array, re-  
20 constructing the failed global diagonal parity storage device using all data and row parity  
21 storage devices of all sub-arrays of the array.

1           12. (Original) The method of Claim 11 wherein the step of correcting storage  
2 device failure further comprises the steps of:

3           if the storage device failure is not a single storage device failure, determining  
4 whether the storage device failure is a double failure within the sub-array;

5           if the storage device failure is not a double failure within the sub-array, determin-  
6 ing whether one of the failures includes the diagonal parity storage device; and  
7 if one of the failures does not include the diagonal parity storage device, reconstructing  
8 the failed storage device in each sub-array using local row parity.

1           13. (Original) The method of Claim 12 wherein the step of correcting storage  
2 device failure further comprises the steps of:

3           if one of the failures includes the diagonal parity storage device, determining  
4 whether another of the failed storage devices includes a row parity storage device;

5           if the another of the failed storage devices includes the row parity storage device,  
6 reconstructing the row parity storage device from the data storage devices of the sub-  
7 array; and

8           reconstructing the diagonal parity storage device from all data and row parity  
9 storage devices of all sub-arrays of the array.

1           14. (Original) The method of Claim 13 wherein the step of correcting storage de-  
2 vice failure further comprises the steps of:

3           if the another of the failed storage devices does not include the row parity storage  
4 device, reconstructing the data storage device using local row parity associated with the  
5 sub-array; and

6           reconstructing the diagonal parity storage device from all data and row parity  
7 storage devices of all sub-arrays of the array.

1           15. (Original) The method of Claim 14 wherein the step of correcting storage  
2 device failures further comprises the step of, if the storage device failure is a double fail-  
3 ure within the sub-array, recovering two failed storage devices within the sub-array using  
4 a row-diagonal reconstruction process.

1           16. (Original) The method of Claim 15 wherein the step of recovering comprises  
2 the steps of:

3           using the diagonal parity storage device to recover at least one data block from a  
4 first of the failed storage devices of the sub-array; and

5           once the data block is recovered, using row parity within the sub-array to recover  
6 a corresponding block in a second of the failed storage devices.

1           17. (Original) The method of Claim 10 further comprising the step of organizing  
2 each sub-array as a concentrated parity disk array.

1           18. (Original) The method of Claim 10 further comprising the step of organizing  
2 each sub-array as a distributed parity disk array.

1           19. (Previously presented) A method for correcting double failures in a storage  
2 array using a combination of a single diagonal parity group and multiple row parity  
3 groups, the method comprising the steps of:

4           organizing the storage array as a plurality of concatenated sub-arrays based on  
5 double failure protection encoding, each sub-array including a set of data storage devices  
6 and a row parity storage device, the storage array further including a global diagonal par-  
7 ity storage device for holding diagonal parity;

8           computing the diagonal parity for the single diagonal parity group across the con-  
9 catenated sub-arrays;

10          correcting storage device failure within the array using the row parity storage de-  
11 vice associated with each sub-array and the global diagonal parity storage device associ-  
12 ated with the storage array;

13          encoding the double failure protection as EVENODD parity encoding;

14          determining whether the storage device failure is to a single storage device in a  
15 sub-array;

16          if the storage device failure is to a single storage device in the sub-array, recon-  
17 structing the failed storage device using local row parity associated with the sub-array;  
18 and

19          if the storage device failure is not to a single storage device in the sub-array, re-  
20 constructing the failed global diagonal parity storage device using all data storage devices  
21 of all sub-arrays of the array.

1           20. (Original) The method of Claim 19 wherein the step of correcting storage  
2 device failure further comprises the steps of:

3           if the storage device failure is not a single storage device failure, determining  
4 whether the storage device failure is a double failure within the sub-array;

5           if the storage device failure is not a double failure within the sub-array, determin-  
6 ing whether one of the failures includes the diagonal parity storage device; and

7 if one of the failures does not include the diagonal parity storage device, reconstructing  
8 the failed storage device in each sub-array using local row parity.

1           21. (Original) The method of Claim 20 wherein the step of correcting storage  
2 device failure further comprises the steps of:

3           if one of the failures includes the diagonal parity storage device, determining  
4 whether another of the failed storage devices includes a row parity storage device;

5           if the another of the failed storage devices includes the row parity storage device,  
6 reconstructing the row parity storage device from the data storage devices of the sub-  
7 array; and

8           reconstructing the diagonal parity storage device from all of the data storage de-  
9 vices of the array.

1           22. (Original) The method of Claim 21 wherein the step of correcting storage  
2 device failure further comprises the steps of:

3           if the another of the failed storage devices does not include the row parity storage  
4 device, reconstructing the data storage device using local row parity associated with the  
5 sub-array; and

6           reconstructing the diagonal parity storage device from all data storage devices of  
7 the array.

1           23. (Original) The method of Claim 22 wherein the step of correcting storage  
2 device failures further comprises the step of, if the storage device failure is a double fail-  
3 ure within the sub-array, recovering two failed storage devices within the sub-array using  
4 an EVENODD reconstruction process.

1           24. (Original) Apparatus for correcting double failures in a storage array using a  
2 combination of a single diagonal parity group and multiple row parity groups, the appa-  
3 rus comprising:  
4           means for organizing the storage array as a plurality of concatenated sub-arrays  
5 based on double failure protection encoding, each sub-array including a set of data stor-  
6 age devices and a row parity storage device, the storage array further including a global  
7 diagonal parity storage device for holding diagonal parity;  
8           means for computing the diagonal parity for the single diagonal parity group  
9 across the concatenated sub-arrays; and  
10          means for correcting storage device failure within the array using the row parity  
11 storage device associated with each sub-array and the global diagonal parity storage de-  
12 vice associated with the storage array.

1           25. (Original) A computer readable medium containing executable program in-  
2 structions for correcting double failures in a storage array using a combination of a single  
3 diagonal parity group and multiple row parity groups, the executable program instruc-  
4 tions comprising program instructions for:  
5           organizing the storage array as a plurality of concatenated sub-arrays based on  
6 double failure protection encoding, each sub-array including a set of data storage devices  
7 and a row parity storage device, the storage array further including a global diagonal par-  
8 ity storage device for holding diagonal parity;



9           computing the diagonal parity for the single diagonal parity group across the con-  
10   catenated sub-arrays;

11           correcting storage device failure within the array using the row parity storage de-  
12   vice associated with each sub-array and the global diagonal parity storage device associ-  
13   ated with the storage array.

1           26. (Original) A system adapted to correct multiple storage element failures in  
2   an array using a combination of multiple first failure recovery groups and a single secon-  
3   dary failure recovery group, the system comprising:

4           a storage array having a plurality of concatenated sub-arrays, each sub-array in-  
5   cluding a set of data storage elements and a first failure recovery storage element storing  
6   first values used to correct a single failure within the sub-array, the array further includ-  
7   ing a global failure recovery storage element associated with the storage array and hold-  
8   ing secondary values for the single secondary failure recovery group, the secondary val-  
9   ues computed across the concatenation of the sub-arrays.

1           27. (Original) The system of Claim 26 wherein the storage elements are packets  
2   and wherein the failure recovery is parity.

1           28. (Previously presented) A method for operating a storage array, comprising:  
2           organizing the storage array as a plurality of concatenated sub-arrays based on  
3   double failure protection encoding, each sub-array including a set of data storage devices  
4   and a row parity storage device, the storage array further including a global diagonal par-  
5   ity storage device for holding diagonal parity;

6           computing the diagonal parity for the single diagonal parity group across the con-  
7   catenated sub-arrays;

8           correcting storage device failure within the array using the row parity storage de-  
9       vice associated with each sub-array and the global diagonal parity storage device associ-  
10      ated with the storage array;

11           determining whether the storage device failure is to a single storage device in a  
12      sub-array;

13           if the storage device failure is to a single storage device in the sub-array, recon-  
14      structing the failed storage device using local row parity associated with the sub-array;  
15      and

16           if the storage device failure is not to a single storage device in the sub-array, re-  
17      constructing the failed global diagonal parity storage device using all data storage devices  
18      of all sub-arrays of the array.

1   29. (Previously presented) A storage array, comprising:

2           means for organizing the storage array as a plurality of concatenated sub-arrays  
3      based on double failure protection encoding, each sub-array including a set of data stor-  
4      age devices and a row parity storage device, the storage array further including a global  
5      diagonal parity storage device for holding diagonal parity;

6           means for computing the diagonal parity for the single diagonal parity group  
7      across the concatenated sub-arrays;

8           means for correcting storage device failure within the array using the row parity  
9      storage device associated with each sub-array and the global diagonal parity storage de-  
10     vice associated with the storage array;

11           means for determining whether the storage device failure is to a single storage de-  
12     vice in a sub-array;

13           if the storage device failure is to a single storage device in the sub-array, means  
14      for reconstructing the failed storage device using local row parity associated with the sub-  
15      array; and

16           if the storage device failure is not to a single storage device in the sub-  
17           array, means for reconstructing the failed global diagonal parity storage device us-  
18           ing all data storage devices of all sub-arrays of the array.

1           30. (Previously presented) A method for correcting double failures in a storage  
2           array, comprising:

3           organizing the storage array as a plurality of concatenated sub-arrays, each sub-  
4           array including a set of data storage devices and a row parity storage device, the storage  
5           array further including a global diagonal parity storage device for holding diagonal par-  
6           ity;

7           computing the diagonal parity across the concatenated sub-arrays; and  
8           correcting storage device failure within the array using the row parity storage de-  
9           vice associated with each sub-array and the global diagonal parity storage device.

1           31. (Previously presented) The method of claim 30, further comprising:  
2           storing all row parity data on a dedicated disk storage device.

1           32. (Previously presented) The method of claim 30, further comprising:  
2           storing all diagonal parity data on a dedicated disk storage device.

1           33. (Previously presented) The method of claim 30, further comprising:  
2           encoding the double failure protection as row-diagonal parity encoding.

1           34. (Previously presented) The method of claim 30, further comprising:  
2           encoding the double failure protection as EVENODD parity encoding.

1           35. (Previously presented) The method of claim 30, further comprising:  
2           determining whether the storage device failure is to a single storage device in a  
3           sub-array;  
4           if the storage device failure is to a single storage device in the sub-array, recon-  
5           structing the failed storage device using local row parity associated with the sub-array;  
6           and  
7           if the storage device failure is not to a single storage device in the sub-array, re-  
8           constructing the failed global diagonal parity storage device using all data and row parity  
9           storage devices of all sub-arrays of the array.

1           36. (Previously presented) The method of claim 30, further comprising:  
2           if the storage device failure is not a single storage device failure, determining  
3           whether the storage device failure is a double failure within the sub-array;  
4           if the storage device failure is not a double failure within the sub-array, determin-  
5           ing whether one of the failures includes the diagonal parity storage device; and  
6           if one of the failures does not include the diagonal parity storage device, recon-  
7           structing the failed storage device in each sub-array using local row parity.

1           37. (Previously presented) The method of claim 30, further comprising:  
2           if one of the failures includes the diagonal parity storage device, determining  
3 whether another of the failed storage devices includes a row parity storage device;  
4           if the another of the failed storage devices includes the row parity storage device,  
5 reconstructing the row parity storage device from the data storage devices of the sub-  
6 array; and  
7           reconstructing the diagonal parity storage device from all data and row parity  
8 storage devices of all sub-arrays of the array.

1           38. (Previously presented) The method of claim 30, further comprising:  
2           if the another of the failed storage devices does not include the row parity storage  
3 device, reconstructing the data storage device using local row parity associated with the  
4 sub-array; and  
5           reconstructing the diagonal parity storage device from all data and row parity  
6 storage devices of all sub-arrays of the array.

1           39. (Previously presented) The method of claim 30, further comprising:  
2           if the storage device failure is a double failure within the sub-array, recovering  
3 two failed storage devices within the sub-array using a row-diagonal reconstruction proc-  
4 ess.

1           40. (Previously presented) The method of claim 30, further comprising:  
2           using the diagonal parity storage device to recover at least one data block from a  
3           first of the failed storage devices of the sub-array; and  
4           once the data block is recovered, using row parity within the sub-array to recover  
5           a corresponding block in a second of the failed storage devices.

1           41. (Previously presented) The method of claim 30, further comprising:  
2           organizing each sub-array as a concentrated parity disk array.

1           42. (Previously presented) The method of claim 30, further comprising:  
2           organizing each sub-array as a distributed parity disk array.

1           43. (Previously presented) A storage array, comprising:  
2           means for organizing the storage array as a plurality of concatenated sub-arrays  
3           based on double failure protection encoding, each sub-array including a set of data stor-  
4           age devices and a row parity storage device, the storage array further including a global  
5           diagonal parity storage device for holding diagonal parity;  
6           means for computing the diagonal parity for the single diagonal parity group  
7           across the concatenated sub-arrays;

8 means for correcting storage device failure within the array using the row parity  
9 storage device associated with each sub-array and the global diagonal parity storage de-  
10 vice associated with the storage array;

11 means for determining whether the storage device failure is to a single storage de-  
12 vice in a sub-array;

13 if the storage device failure is to a single storage device in the sub-array, means  
14 for reconstructing the failed storage device using local row parity associated with the sub-  
15 array; and

16 if the storage device failure is not to a single storage device in the sub-array,  
17 means for reconstructing the failed global diagonal parity storage device using all data  
18 storage devices of all sub-arrays of the array.

1 44. (Previously presented) A storage array, comprising:

2 means for organizing the storage array as a plurality of concatenated sub-arrays  
3 based on double failure protection encoding, each sub-array including a set of data stor-  
4 age devices and a row parity storage device, the storage array further including a global  
5 diagonal parity storage device for holding diagonal parity;

6 means for computing the diagonal parity for the single diagonal parity group  
7 across the concatenated sub-arrays;

8 means for correcting storage device failure within the array using the row parity  
9 storage device associated with each sub-array and the global diagonal parity storage de-  
10 vice associated with the storage array;

11 means for determining whether the storage device failure is to a single storage de-  
12 vice in a sub-array;

13 if the storage device failure is to a single storage device in the sub-array, means  
14 for reconstructing the failed storage device using local row parity associated with the sub-  
15 array; and

16           if the storage device failure is not to a single storage device in the sub-array,  
17   means for reconstructing the failed global diagonal parity storage device using all data  
18   storage devices of all sub-arrays of the array.

1           45. (Previously presented) A storage array, comprising:

2           means for organizing the storage array as a plurality of concatenated sub-arrays,  
3   each sub-array including a set of data storage devices and a row parity storage device, the  
4   storage array further including a global diagonal parity storage device for holding diago-  
5   nal parity;

6           means for computing the diagonal parity across the concatenated sub-arrays; and

7           means for correcting storage device failure within the array using the row parity  
8   storage device associated with each sub-array and the global diagonal parity storage de-  
9   vice.

1           46. (Previously presented) The method of claim 45, further comprising:

2           means for storing all row parity data on a dedicated disk storage device.

1           47. (Previously presented) The method of claim 45, further comprising:

2           means for storing all diagonal parity data on a dedicated disk storage de-  
3   vice.

1           48. (Previously presented) The storage array of claim 45, further comprising:

2           means for encoding the double failure protection as row-diagonal parity encoding.



1           49. (Previously presented) The storage array of claim 45, further comprising:  
2           means for encoding the double failure protection as EVENODD parity encoding.

1           50. (Previously presented) The storage array of claim 45, further comprising:  
2           means for determining whether the storage device failure is to a single storage de-  
3           vice in a sub-array;  
4           if the storage device failure is to a single storage device in the sub-array, means  
5           for reconstructing the failed storage device using local row parity associated with the sub-  
6           array; and  
7           if the storage device failure is not to a single storage device in the sub-array,  
8           means for re-constructing the failed global diagonal parity storage device using all data  
9           and row parity storage devices of all sub-arrays of the array.

1           51. (Previously presented) The storage array of claim 45, further comprising:  
2           if the storage device failure is not a single storage device failure, means for de-  
3           termining whether the storage device failure is a double failure within the sub-array;  
4           if the storage device failure is not a double failure within the sub-array, means for  
5           determining whether one of the failures includes the diagonal parity storage device; and  
6           if one of the failures does not include the diagonal parity storage device, means  
7           for reconstructing the failed storage device in each sub-array using local row parity.

1           52. (Previously presented) The storage array of claim 45, further comprising:  
2           if one of the failures includes the diagonal parity storage device, means for deter-  
3 mining whether another of the failed storage devices includes a row parity storage device;  
4           if the another of the failed storage devices includes the row parity storage device,  
5 means for reconstructing the row parity storage device from the data storage devices of  
6 the sub-array; and  
7           means for reconstructing the diagonal parity storage device from all data and row  
8 parity storage devices of all sub-arrays of the array.

1           53. (Previously presented) The storage array of claim 45, further comprising:  
2           if the another of the failed storage devices does not include the row parity storage  
3 device, means for reconstructing the data storage device using local row parity associated  
4 with the sub-array; and  
5           means for reconstructing the diagonal parity storage device from all data and row  
6 parity storage devices of all sub-arrays of the array.

1           54. (Previously presented) The storage array of claim 45, further comprising:  
2           if the another of the failed storage devices does not include the row parity storage  
3 device, means for reconstructing the data storage device using local row parity associated  
4 with the sub-array; and

5 means for reconstructing the diagonal parity storage device from all data and row  
6 parity storage devices of all sub-arrays of the array.

1 55. (Previously presented) The storage array of claim 45, further comprising:

2 means for using the diagonal parity storage device to recover at least one data  
3 block from a first of the failed storage devices of the sub-array; and

4 once the data block is recovered, means for using row parity within the sub-array  
5 to recover a corresponding block in a second of the failed storage devices.

1 56. (Previously presented) The storage array of claim 45, further comprising:

2 means for organizing each sub-array as a concentrated parity disk array.

1 57. (Previously presented) The storage array of claim 45, further comprising:

2 means for organizing each sub-array as a distributed parity disk array.

1 58. (Previously presented) A computer readable media, comprising:

2 said computer readable media containing instructions for execution on a processor  
3 for the practice of a method for correcting double failures in a storage array, having the  
4 steps,

5 organizing the storage array as a plurality of concatenated sub-arrays, each sub-  
6 array including a set of data storage devices and a row parity storage device, the storage

7 array further including a global diagonal parity storage device for holding diagonal par-  
8 ity;  
9 computing the diagonal parity across the concatenated sub-arrays; and  
10 correcting storage device failure within the array using the row parity storage de-  
11 vice associated with each sub-array and the global diagonal parity storage device.

1 59. (Previously presented) Electromagnetic signals propagating on a computer  
2 network, comprising:

3 said electromagnetic signals carrying instructions for execution on a processor for  
4 the practice of a method for correcting double failures in a storage array, having the steps,  
5 organizing the storage array as a plurality of concatenated sub-arrays, each sub-  
6 array including a set of data storage devices and a row parity storage device, the storage  
7 array further including a global diagonal parity storage device for holding diagonal par-  
8 ity;  
9 computing the diagonal parity across the concatenated sub-arrays; and  
10 correcting storage device failure within the array using the row parity storage de-  
11 vice associated with each sub-array and the global diagonal parity storage device.